

Parameter Compilation

Hubie Chen *

March 3, 2015

In resolving instances of a computational problem, if multiple instances of interest share a feature in common, it may be fruitful to compile this feature into a format that allows for more efficient resolution, even if the compilation is relatively expensive. In this article, we introduce a formal framework for classifying problems according to their compilability. The basic object in our framework is that of a parameterized problem, which here is a language along with a parameterization—a map which provides, for each instance, a so-called parameter on which compilation may be performed. Our framework is positioned within the paradigm of parameterized complexity, and our notions are relatable to established concepts in the theory of parameterized complexity. Indeed, we view our framework as playing a unifying role, integrating together parameterized complexity and compilability theory.

1. Introduction

In resolving instances of a computational problem, if it is the case that multiple instances of interest share a feature in common, it may be fruitful to *compile* this feature into a format that allows for more efficient resolution, even if the compilation is relatively expensive. As a first, simple example, consider the problem of deciding if two nodes of an undirected graph are connected. If it is anticipated that many such connectivity queries will share the same graph G , it may be worthwhile to compile G into a format that will allow for accelerated resolution of the queries. As a second example, consider the problem of evaluating a database query on a database. If one is interested in a small set of queries that will be posed to numerous databases, it may be worthwhile to compile the queries of interest into a format that allows for the fastest evaluation. Note that a relatively expensive compilation process may be worthwhile if its results are amortized by repeated use. Indeed, one may conceive of compilation as an off-line preprocessing, whose expense is offset by its later on-line use.

*Universidad del País Vasco, E-20018 San Sebastián, Spain, *and* IKERBASQUE, Basque Foundation for Science, E-48011 Bilbao, Spain

In this article, we attempt to make an infrastructural contribution by introducing a formal framework for classifying problems according to their compilability. Such a framework was previously presented by Cadoli, Domini, Liberatore, and Schaerf [2], (hereafter, *CDLS*); we will discuss the relationship between our framework and theirs below.

The basic object in our framework is a *parameterized problem*, which we define to be a language Q along with a *parameterization* κ , a polynomial-time computable mapping defined from strings to strings. (For precise details and justifications of definitions, refer to the technical sections of the article.) As usual, we refer to $\kappa(x)$ as the *parameter* of an instance x . In our framework, we wish to understand for which problems the parameters can be succinctly compiled into a form such that, post-compilation, the problem can be resolved in polynomial-time. The base class of our framework, called **poly-comp-PTIME**, is (essentially) defined to contain a parameterized problem (Q, κ) if there exists a polynomial-length, computable function c such that if each instance x is always presented along with $c(\kappa(x))$, then each instance can be resolved in polynomial time. The function c models the notion of compilation of the parameters. In order to give evidence of non-containment in the class **poly-comp-PTIME** and also to facilitate problem classification, we introduce a hierarchy of parameterized complexity classes **chopped- \mathcal{C}** , one for each classical complexity class \mathcal{C} ; we observe (for example) that **chopped-NP** is not contained in **poly-comp-PTIME**, assuming that the polynomial hierarchy does not collapse (see Proposition 9 and Theorem 19), and hence hardness of a problem for **chopped-NP** can be construed as evidence of non-containment in **poly-comp-PTIME**. We observe a number of completeness and hardness results for **chopped-NP** (Section 6).¹ The class **poly-comp-PTIME** and the classes **chopped- \mathcal{C}** are all subsets of the parameterized class **FPT**, which is considered to be the basic notion of tractability in the paradigm of parameterized complexity.² We believe that the introduced classes constitute a natural stratification of **FPT**, whose study might well lead to deeper theory.

In the *CDLS* framework, the basic object is a language consisting of pairs of strings (called a *language of pairs*), and one aims to understand when a compilation can be applied to the first entry of each pair so as to allow for efficient decision. This is a point of difference with our framework, but note that the notions from our framework can be readily applied to the languages of pairs that *CDLS* study by using the parameterization π_1 that returns the first entry of a pair. Another point of difference between our framework and theirs is that their analog of our compilation function c is not required to be computable; while this makes the negative results stronger, in our view there ought to be a focus on positive results, which are rendered less meaningful without the computability requirement. (Actually, we are not aware of any natural computable problem for which the presence or absence of this requirement makes a difference.) Although these differences may appear slight, by initiating our theory with our particular choice

¹ These results include a hardness result on model checking existential positive sentences (Proposition 31); we remark that obtaining a broader understanding of the non-compilability results in the author's previous study of model checking [5] was in fact a motivation of the present article.

² Note that the containment of **poly-comp-PTIME** in **FPT** is essentially observed (in different language) in the last paragraph of Section 5 of [2].

of definitions, we are able to position our framework within the language and tradition of parameterized complexity and relate our notions to existing ideas in parameterized complexity. For instance, although not difficult, we can directly relate the notion of a *polynomial kernelization* to the classes **chopped- \mathcal{C}** (Proposition 10) and use this relationship to observe the **chopped-NP**-completeness of the standard parameterization of the hitting set problem for hypergraphs of bounded edge size (see Theorem 30). We also believe that the theory that results from our framework’s definitions witnesses that working with parameterized problems as opposed to languages of pairs allows for greater flexibility and smoother formulation (consider, for example, the characterization of **chopped- \mathcal{C}** using the *length parameterization* given by Proposition 14).

Our framework and that of CDLS also differ later in the respective developments. Notably, our notion of reduction (Definition 11) is readily seen to be a restricted version of the usual fpt many-one reduction in parameterized complexity, and we believe that our notion of reduction is conceptually simpler to comprehend than that of CDLS [2, Definition 2.8]. Despite these differences—and we view this as crucial—we demonstrate how classification results obtained in the CDLS framework can be formulated and obtained in our framework; this is made precise and performed in Section 5.

The presentation and development of our framework may thus be viewed as playing a unifying role, integrating together parameterized complexity and compilation. Our choices of definitions and in formulation allows us to directly relate the resulting concepts to the theory of parameterized complexity. At the same time, we believe that these concepts capture in an essential way the core mathematical content and the core ideas of the CDLS framework (as borne out by our results and discussion in Section 5).

Related work. The CDLS framework was deployed after its introduction to analyze the compilability of reasoning tasks, see for example [10, 11].

In the context of compiling propositional formulas, a notion of compilation whereby a compiled version should have the same models as the original formula was studied, for example by Gogic et al. [9] and by Darwiche and Marquis [6]; see also the recent work by Bova et al. [1].

Variants of the CDLS framework that relaxed the requirement that the size of compilations be polynomial were also studied [3, 4].

Finally, we mention that Fan, Geerts, and Neven [7] also developed a framework for classifying problems according to compilability, with a focus on efficient parallel processing (modelled using the complexity class NC) following a polynomial-time compilation. We believe that it may be of interest to better understand and develop the relationship between our framework and theirs. While we leave such a study to future work, we mention that their notion of Π -tractability on a language Q of pairs can be described using our framework.³

³ Precisely, a language Q of pairs being Π -tractable can be verified to be equivalent to the parameterized problem (Q, π_1) being in our class **poly-comp-NC** via a poly-compilable function $g(x, y) = f(c(\pi_1(x, y)), (x, y)) = f(c(x), (x, y))$ where c is polynomial-time computable.

2. Preliminaries

Throughout, π_i denotes the operator that, given a tuple, returns the i th entry of the tuple.

When T is a set, we use $\wp_{\text{fin}}(T)$ to denote the set $\{S \subseteq T \mid S \text{ is finite}\}$.

We generally use Σ to denote the alphabet over which strings are formed, and generally assume $\{0, 1\} \subseteq \Sigma$. As is standard, we freely interchange between elements of Σ^* and $\Sigma^* \times \Sigma^*$. When $k \geq 0$, we use $\Sigma^{\leq k}$ to denote the set of strings in Σ^* of length less than or equal to k . For $n \in \mathbb{N}$, we use $\text{un}(n)$ to denote its unary encoding 1^n as a string.

We assume that languages under discussion are non-trivial, that is, not equal to \emptyset nor Σ^* . We use **PTIME** to denote the set of all languages decidable in polynomial time, and **fPTIME** to denote the set of all functions from Σ^* to Σ^* that are computable in polynomial time.

Here, by a *parameterization*, we refer to a map from Σ^* to Σ^* . Relative to a parameterization $\kappa : \Sigma^* \rightarrow \Sigma^*$, it is typical to refer to $\kappa(x)$ as the *parameter* of the string x . While it is typical in the literature to define a parameterization to be a map from Σ^* to \mathbb{N} , in this article we want to apply compilation functions to parameters and discuss the *length* of the results, and we find that this is facilitated in many cases by permitting the parameter of a string to be a string itself. Throughout, we employ the following assumption (which is discussed below in Remark 4).

Assumption 1. *Each parameterization is polynomial-time computable, that is, in **fPTIME**.*

We use len to denote the parameterization defined by $\text{len}(x) = \text{un}(|x|)$. A *parameterized problem* is a pair (Q, κ) consisting of a language Q and a parameterization κ .

By a *classical complexity class*, we refer to a set of computable languages. For a classical complexity class \mathcal{C} , we define **para- \mathcal{C}** to be the set that contains a parameterized problem (Q, κ) if there exists a computable function $c : \Sigma^* \rightarrow \Sigma^*$, and a language $Q' \subseteq \Sigma^* \times \Sigma^*$ in \mathcal{C} such that, for each string $x \in \Sigma^*$, it holds that $x \in Q \Leftrightarrow (c(\kappa(x)), x) \in Q'$. We define **FPT** to be **para-PTIME** (although this is perhaps not the usual definition of **FPT**, it is equivalent [8, Theorem 1.37]).

As usual, when \mathcal{D} is a set of problems (that is, a set of either languages or parameterized problems), we say that a problem P' is *\mathcal{D} -hard* under a notion of reduction if each P in \mathcal{D} reduces to P' ; if in addition $P' \in \mathcal{D}$, we say that P' is *\mathcal{D} -complete*. We say that \mathcal{D} is *closed* under a notion of reduction if, when P reduces to P' and $P' \in \mathcal{D}$, it holds that $P \in \mathcal{D}$.

3. Framework

3.1. Problem classes

In this subsection, we introduce the complexity classes of our framework. We begin by introducing two basic definitions. By a *length function*, we refer to a function from \mathbb{N} to \mathbb{N} .

Definition 2. Let \mathcal{L} be a set of length functions.

- A function $c : \Sigma^* \rightarrow \Sigma^*$ is said to be \mathcal{L} -length if there exists $\ell \in \mathcal{L}$ such that for each $x \in \Sigma^*$, it holds that $|c(x)| \leq \ell(|x|)$.
- A function $g : \Sigma^* \rightarrow \Sigma^*$ is \mathcal{L} -*compilable* with respect to a parameterization κ if there exist $f \in \text{fPTIME}$ and a computable, \mathcal{L} -length function $c : \Sigma^* \rightarrow \Sigma^*$ such that (for each $x \in \Sigma^*$) $g(x) = f(c(\kappa(x)), x)$.

Put informally, a function g is \mathcal{L} -compilable if, when one has the result of applying c to the parameter of an instance x , the value $g(x)$ can be efficiently computed. The function c can be thought of as performing a precomputation or compilation of the parameter. Here, we do not place any restriction on the computational resources needed to compute c , other than requiring that c is computable. We view the requirement that c be computable as natural in terms of claiming positive results, as we find it hard to argue that a non-computable compilation would actually be usable. We do restrict the length of c according to \mathcal{L} ; we will be most interested in the case where the length of c is polynomially bounded.

With this terminology in hand, we can now define our first classes of parameterized problems.

Definition 3. Let \mathcal{L} be a set of length functions, and let \mathcal{C} be a classical complexity class.

- We say that a parameterized problem (Q, κ) is \mathcal{L} -*compilable to* \mathcal{C} if there exists a function $g : \Sigma^* \rightarrow \Sigma^*$ that is \mathcal{L} -compilable (with respect to κ) and a language $Q' \in \mathcal{C}$ such that (for each $x \in \Sigma^*$) $x \in Q \Leftrightarrow g(x) \in Q'$.
- We define $\mathcal{L}\text{-comp-}\mathcal{C}$ to be the set that contains each parameterized problem that is \mathcal{L} -compilable to \mathcal{C} .

When \mathcal{L} is the set of all polynomials on \mathbb{N} , we define $\text{poly-comp-}\mathcal{C}$ as $\mathcal{L}\text{-comp-}\mathcal{C}$ and speak, for instance, of *poly-compilability*; similarly, when \mathcal{L} is the set $\cup\{O(2^p) \mid p \text{ is a polynomial}\}$ of *exponential functions*, we define $\text{exp-comp-}\mathcal{C}$ as $\mathcal{L}\text{-comp-}\mathcal{C}$ and speak, for instance, of *exp-compilability*.

Remark 4. In this paper, the smallest class that we will consider is poly-comp-PTIME , and we will regard an inclusion result in this class as the most positive result demonstrable on a parameterized problem. Suppose that a parameterized problem (Q, κ) is in poly-comp-PTIME via $g(x) = f(c(\kappa(x)), x)$ and Q' . One way to intuitively interpret this inclusion is as follows. Suppose that the value $c(k)$ is known for parameter values k in a limited range. Then, for each instance $x \in \Sigma^*$ having parameter value $\kappa(x)$ in that limited range, whether or not $x \in Q$ can be determined efficiently, by applying the efficiently computable function f to $(c(\kappa(x)), x)$ and then by invoking an efficient decision procedure for Q' . Indeed, our intention here is to model the notion of efficient decidability modulo knowledge of c ; this is why we put into effect Assumption 1.

We observe the following upper bound on each class $\mathcal{L}\text{-comp-}\mathcal{C}$, which in particular indicates that $\text{poly-comp-PTIME} \subseteq \text{FPT}$.

Proposition 5. *Let \mathcal{L} be a set of length functions, and let \mathcal{C} be a classical complexity class that is closed under many-one polynomial-time reduction. It holds that $\mathcal{L}\text{-comp-}\mathcal{C} \subseteq \text{para-}\mathcal{C}$.*

Proof. Suppose that (Q, κ) is \mathcal{L} -compilable to \mathcal{C} via $g(x) = f(c(\kappa(x)), x)$ and $Q'' \in \mathcal{C}$, so that $x \in Q \Leftrightarrow g(x) \in Q''$. Define $Q' = \{(a, b) \mid f(a, b) \in Q''\}$. The language Q' is many-one polynomial-time reducible to Q'' via f , so $Q' \in \mathcal{C}$. We have $x \in Q \Leftrightarrow (c(\kappa(x)), x) \in Q'$, implying that $Q \in \text{para-}\mathcal{C}$. \square

We now define a family of complexity classes which will be used to classify parameterized problems in FPT according to their compilability, and in particular to give evidence of non-inclusion in poly-comp-PTIME , via hardness results.

Definition 6. *For each classical complexity class \mathcal{C} , we define $\text{chopped-}\mathcal{C}$ to be the set that contains each parameterized problem (Q, κ) that is in $\text{poly-comp-}\mathcal{C}$ via a function g for which there exists a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that (for each $x \in \Sigma^*$) $|g(x)| \leq p(|\kappa(x)|)$.*

For the sake of understanding this definition, let us call the restriction of a language Q' to $Q' \cap \Sigma^{\leq k}$ the *chop having magnitude k* of Q' . Then, intuitively speaking, a problem is in $\text{chopped-}\mathcal{C}$ if it is in $\text{poly-comp-}\mathcal{C}$ via g and Q' where $g(x)$ accesses only a chop (of Q') having magnitude restricted by a polynomial in the parameter of x . The following proposition is clear from the definition of $\text{chopped-}\mathcal{C}$.

Proposition 7. *For each classical complexity class \mathcal{C} , it holds that*

$$\text{chopped-}\mathcal{C} \subseteq \text{poly-comp-}\mathcal{C}.$$

We also have the following upper bound on $\text{chopped-}\mathcal{C}$, which shows that the classes $\text{chopped-}\mathcal{C}$ constitute a stratification of the class FPT .

Proposition 8. *For each classical complexity class \mathcal{C} , it holds that*

$$\text{chopped-}\mathcal{C} \subseteq \text{exp-comp-PTIME},$$

and hence that $\text{chopped-}\mathcal{C} \subseteq \text{FPT}$ (by Proposition 5).

Proof. We prove that $\text{chopped-}\mathcal{C} \subseteq \text{exp-comp-PTIME}$. Fix $x_N, x_Y \in \Sigma^*$ and $P \in \text{PTIME}$ such that $x_Y \in P$ and $x_N \notin P$. Assume that (Q, κ) is in $\text{chopped-}\mathcal{C}$ via $g(x) = f(c(\kappa(x)), x)$, the polynomial p , and $Q' \in \mathcal{C}$. Let $c_1^+ : \Sigma^* \rightarrow \Sigma^*$ be the function computed by the algorithm that, given $k \in \Sigma^*$, loops over each string y in $\Sigma^{\leq p(|k|)}$ and, for each such string y , outputs 1 or 0 depending on whether or not $y \in Q'$; thus, $|c_1^+(k)| = |\Sigma^{\leq p(|k|)}|$. Define $c^+(k) = (c_1^+(k), c(k), k)$. Let f^+ be a function computed by a polynomial-time algorithm that, given a string $((d_1, d, k), x)$ where d_1 is a string over $\{0, 1\}$ of length $|\Sigma^{\leq p(|k|)}|$, computes $x' = f(d, x)$, computes the bit b of d_1 corresponding to x' (whenever

$|x'| \leq p(|k|)$, and outputs x_Y or x_N depending on whether or not $b = 1$ or $b = 0$. The function $g^+(x) = f^+(c^+(\kappa(x)), x)$ witnesses that (Q, κ) is exp-compilable to PTIME: We have that $x \in Q$ iff $f(c(\kappa(x)), x) \in Q'$ iff $f^+((c_1^+(\kappa(x), c(\kappa(x))), \kappa(x)), x) = x_Y$ iff $f^+(c^+(\kappa(x)), x) \in P$. \square

We observe that our base class **poly-comp-PTIME** coincides with the class **chopped-PTIME**, which is the smallest class that we will consider from the hierarchy of classes **chopped-C**.

Proposition 9. **chopped-PTIME** = **poly-comp-PTIME**.

Proof. The \subseteq direction follows from Proposition 7. For the \supseteq direction, suppose that $(Q, \kappa) \in \text{poly-comp-PTIME}$. Then, there exists a function g that is a poly-compilable with respect to κ and $Q' \in \text{PTIME}$ such that $x \in Q \Leftrightarrow g(x) \in Q'$. Fix $x_N, x_Y \in \Sigma^*$ and $P \in \text{PTIME}$ such that $x_Y \in P$ and $x_N \notin P$. Let $h \in \text{fPTIME}$ be a function such that, for all $x \in \Sigma^*$, it holds that $x \in Q'$ implies $h(x) = x_Y$ and that $x \notin Q'$ implies $h(x) = x_N$. Then, the mapping $h(g)$ witnesses that (Q, κ) is in **chopped-PTIME**:

$$x \in Q \Leftrightarrow g(x) \in Q' \Leftrightarrow h(g(x)) \in P$$

and for all strings $x \in \Sigma^*$, it holds that $|h(g(x))| \leq \max(|x_Y|, |x_N|)$. \square

The classes **chopped-C** can be directly related to kernelization in the following way. Here, we say that a parameterized problem (Q, κ) has a *polynomial kernelization* if there exists a polynomial-time computable function $K : \Sigma^* \rightarrow \Sigma^*$ and a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that (for each $x \in \Sigma^*$) $x \in Q \Leftrightarrow K(x) \in Q$ and $|K(x)| \leq p(|\kappa(x)|)$.

Proposition 10. *Suppose that a parameterized problem (Q, κ) has a polynomial kernelization and \mathcal{C} is a classical complexity class such that $Q \in \mathcal{C}$. Then, the problem (Q, κ) is in **chopped-C**.*

Proof. We have that (Q, κ) in **poly-comp-C** via K (the function from the definition of polynomial kernelization), since $x \in Q \Leftrightarrow K(x) \in Q$. Moreover, it holds that there exists a polynomial p such that $|K(x)| \leq p(|\kappa(x)|)$ by the definition of polynomial kernelization. \square

3.2. Reduction

We now introduce a notion of reduction for comparing the compilability of parameterized problems.

Definition 11. *We say that a parameterized problem (Q, κ) poly-comp reduces to another parameterized problem (Q', κ') if there exists a function $g : \Sigma^* \rightarrow \Sigma^*$ that is poly-compilable with respect to κ and a poly-length, computable function $s : \Sigma^* \rightarrow \wp_{\text{fin}}(\Sigma^*)$ such that (for each $x \in \Sigma^*$) it holds that $x \in Q \Leftrightarrow g(x) \in Q'$ and that $\kappa'(g(x)) \in s(\kappa(x))$.*

The notion of poly-comp reduction can be viewed as a restricted version of fpt many-one reduction. (Consider, for example, the definition given by Flum and Grohe [8,

Definition 2.1]; the function g in Definition 11 can be seen to be computable by a fpt-algorithm, and the condition on the function s ensures that their condition (3), when reformulated for parameterizations of the type considered here, holds.)

Note that, in Definition 11, we assume that the set $s(x)$ is given according to a standard representation that lists the strings therein; hence, as a consequence of the assumption that s is poly-length, the size $|s(x)|$ of $s(x)$ is bounded above by a polynomial in $|x|$.

We have the following two basic properties of poly-comp reduction.

Theorem 12. *For each classical complexity class \mathcal{C} , it holds that $\text{poly-comp-}\mathcal{C}$ is closed under poly-comp reduction.*

Theorem 13. *Poly-comp reducibility is transitive.*

We establish these theorems in the appendix (Section A).

We now give an alternative characterization of $\text{chopped-}\mathcal{C}$ in terms of poly-comp reduction.

Proposition 14. *Let \mathcal{C} be a classical complexity class. A parameterized problem (Q, κ) is in $\text{chopped-}\mathcal{C}$ if and only if there exists a language $Q' \in \mathcal{C}$ such that (Q, κ) poly-comp reduces to (Q', len) .*

Proof. For the forward direction, suppose $(Q, \kappa) \in \text{chopped-}\mathcal{C}$. There exists a function g that is poly-compilable with respect to κ , a language $Q' \in \mathcal{C}$, and a polynomial p such that: for each $x \in \Sigma^*$, it holds that $x \in Q$ iff $g(x) \in Q'$ and it holds that $|g(x)| \leq p(|\kappa(x)|)$. Define $s(x) = \{\text{un}(0), \text{un}(1), \dots, \text{un}(p(|\kappa(x)|))\}$; then, for each $x \in \Sigma^*$, it holds that $\kappa'(g(x)) = \text{un}(|g(x)|) \in s(x)$. So, (g, s) is a poly-comp reduction from (Q, κ) to (Q', len) .

For the backward direction, suppose that (Q, κ) poly-comp reduces to (Q', len) via (g, s) , where $Q' \in \mathcal{C}$. We have, for each $x \in \Sigma^*$ that $x \in Q \Leftrightarrow g(x) \in Q'$ and that $\text{len}(g(x)) \in s(\kappa(x))$. But there exists a polynomial p such that for each $x \in \Sigma^*$ and each $y \in s(\kappa(x))$, it holds that $|y| \leq p(|x|)$, since κ and s are both poly-length. Thus $\text{len}(g(x)) = \text{un}(|g(x)|) \in s(\kappa(x))$ implies that $|g(x)| = |\text{un}(|g(x)|)| \leq p(|x|) = p(|\text{len}(x)|)$. \square

From the just-given characterization of $\text{chopped-}\mathcal{C}$, we may infer the following two results.

Proposition 15. *For each classical complexity class \mathcal{C} , the class $\text{chopped-}\mathcal{C}$ is closed under poly-comp reduction.*

Proof. This is a consequence of Proposition 14 and Theorem 13. \square

When discussing a class $\text{chopped-}\mathcal{C}$, we assume by default that hardness and completeness are with respect to poly-comp reducibility.

Proposition 16. *Let \mathcal{C} be a classical complexity class and assume that Q^+ is \mathcal{C} -complete under many-one polynomial-time reduction. Then, the parameterized problem (Q^+, len) is complete for $\text{chopped-}\mathcal{C}$.*

Proof. The problem (Q^+, len) is in $\text{chopped-}\mathcal{C}$ by Proposition 14. Let (Q, κ) be an arbitrary problem in $\text{chopped-}\mathcal{C}$. By Proposition 14, there exists a language $Q' \in \mathcal{C}$ such that (Q, κ) poly-comp reduces to (Q', len) . By hypothesis, there exists a many-one polynomial-time reduction from Q' to Q^+ ; it is straightforward to verify that this implies that (Q', len) poly-comp reduces to (Q^+, len) . The result then follows from Theorem 13. \square

4. Chopped classes and advice

In this section, we relate the classes $\text{chopped-}\mathcal{C}$ to advice-based complexity classes; this will allow us to provide evidence of separation between classes of the form $\text{chopped-}\mathcal{C}$.

We first present a known notion from computational complexity theory, the notion of an advice version of a complexity class. For each classical complexity class \mathcal{C} , we define \mathcal{C}/poly to be the set that contains a language Q if and only if there exists a poly-length map $a : \{1\}^* \rightarrow \Sigma^*$ and a language $Q' \in \mathcal{C}$ such that, for each $x \in \Sigma^*$, it holds that $x \in Q \Leftrightarrow (a(\text{un}(|x|)), x) \in Q'$.

The following theorem shows that containment of one chopped class in another implies a containment in classical complexity.

Theorem 17. *Let \mathcal{C} and \mathcal{C}' be classical complexity classes where \mathcal{C}' is closed under many-one polynomial-time reduction. If $\text{chopped-}\mathcal{C} \subseteq \text{chopped-}\mathcal{C}'$, then $\mathcal{C} \subseteq \mathcal{C}'/\text{poly}$.*

To prove this theorem, we first establish a lemma.

Lemma 18. *Let \mathcal{C}' be a classical complexity class that is closed under many-one polynomial-time reduction. If Q is a language such that $(Q, \text{len}) \in \text{chopped-}\mathcal{C}'$, then $Q \in \mathcal{C}'/\text{poly}$.*

Proof. By hypothesis, there exists a language $Q' \in \mathcal{C}'$ and a function $g : \Sigma^* \rightarrow \Sigma^*$ that is poly-compilable with respect to len such that $x \in Q \Leftrightarrow g(x) \in Q'$. Let us denote $g(x) = f(c(\text{len}(x)), x)$. Let P be the set of pairs $\{(d, x) \in \Sigma^* \times \Sigma^* \mid f(d, x) \in Q'\}$. Then P many-one polynomial-time reduces to Q' via f , so by hypothesis $P \in \mathcal{C}$. Define $a : \{1\}^* \rightarrow \Sigma^*$ by $a(1^n) = c(1^n)$. We have $Q \in \mathcal{C}'/\text{poly}$ via a and Q' , as $x \in Q$ iff $f(a(\text{un}(|x|)), x) = f(c(\text{len}(x)), x) \in Q'$ iff $(a(\text{un}(|x|)), x) \in P$. \square

Proof. (Theorem 17) Suppose that $Q \in \mathcal{C}$. By Proposition 14, it holds that (Q, len) is in $\text{chopped-}\mathcal{C}$. By hypothesis, it holds that (Q, len) is in $\text{chopped-}\mathcal{C}'$. By Lemma 18, it follows that $Q \in \mathcal{C}'/\text{poly}$. \square

We use Σ_i^{P} and Π_i^{P} (with $i \geq 0$) to denote the classes of the polynomial hierarchy (PH); recall that $\Sigma_0^{\text{P}} = \Pi_0^{\text{P}} = \text{PTIME}$, $\Sigma_1^{\text{P}} = \text{NP}$, and $\Pi_1^{\text{P}} = \text{co-NP}$. For each $i \geq 0$, let us say that the classes Σ_i and Π_i are at the i th level of the PH. Let us say that a class \mathcal{C}' of the PH is *above* another class \mathcal{C} of the PH if they are equal or if the level j of \mathcal{C}' is strictly greater than the level i of \mathcal{C} .

Theorem 19. *(follows from [13]) Suppose that \mathcal{C} and \mathcal{C}' are classes of the PH such that \mathcal{C}' is not above \mathcal{C} .*

- If $\text{chopped-}\mathcal{C} \subseteq \text{chopped-}\mathcal{C}'$, then the PH collapses.
- A parameterized problem (Q, κ) that is $\text{chopped-}\mathcal{C}$ -hard is not in $\text{chopped-}\mathcal{C}'$, unless the PH collapses.

Proof. For the first claim, it follows from Theorem 17 that $\mathcal{C} \subseteq \mathcal{C}'/\text{poly}$; by [13], it follows that the PH collapses. For the second claim, observe that if (Q, κ) is $\text{chopped-}\mathcal{C}$ -hard, then $(Q, \kappa) \in \text{chopped-}\mathcal{C}'$ implies that $\text{chopped-}\mathcal{C} \subseteq \text{chopped-}\mathcal{C}'$, by the closure of $\text{chopped-}\mathcal{C}'$ under poly-comp reduction (Theorem 12). \square

5. Relationship to the CDLS framework

In this section, we discuss the relationship between our framework and the CDLS framework. We in particular show that, in a sense that we make precise, the completeness results that they obtain for their problem classes can be formulated and obtained in our framework. Note that the proofs of Theorems 21, 25, and 28 are deferred to the appendix (Section B).

By a *language of pairs*, we refer to a subset of $\Sigma^* \times \Sigma^*$.

The CDLS framework defines, for each classical complexity class, a class which they refer to as the *class of problems non-uniformly compilable to a class C* , and which contains languages of pairs [2, Definition 2.7]. We give the following formulation of this definition.

Definition 20. A language $B \subseteq \Sigma^* \times \Sigma^*$ of pairs is in *mixed- \mathcal{C}* if there exists a polynomial-length, computable function $f : \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$ and a language $B' \in \mathcal{C}$ of pairs such that $(x, y) \in B \Leftrightarrow (f(x, \text{un}(|y|)), y) \in B'$.

Note that our definition is not exactly equivalent to theirs; we require that the function f is computable, while they do not. We do not know of any natural language of pairs for which this makes a difference; assuming computability of f will allow us to more readily relate the defined classes to those of our framework.

To illustrate how classification results on languages obtained in the CDLS framework can be obtained in our framework, we discuss three running examples (studied in [2]):

- Define CI (*clause inference*) to be the set of pairs (x, y) where x is a propositional 3CNF formula, y is a clause, and $x \models y$. We assume here that clauses do not contain repeated literals.
- Define MMC (*minimal model checking*) to be the set of pairs (x, y) where x is a propositional formula and y is a minimal model of x . By *minimal*, we mean with respect to the order \leq where $z \leq z'$ if and only if all variables true under z are also true under z' .
- Define CMI (*clause minimal inference*) to be the set of pairs (x, y) where x is a propositional formula and y is a clause that is satisfied by all minimal models of x .

It is known and straightforward to verify that $\text{CI}, \text{MMC} \in \text{co-NP}$ and $\text{CMI} \in \Pi_2^P$. It follows immediately that $\text{CI}, \text{MMC} \in \text{mixed-co-NP}$ and $\text{CMI} \in \text{mixed-}\Pi_2^P$.

Let us say that a parameterized problem (Q, κ) has *poly-bounded slices* if there exists a polynomial p such that, for each $x \in Q$, it holds that $|x| \leq p(|\kappa(x)|)$. Each of the three parameterized problems (CI, π_1) , (MMC, π_1) , and (CMI, π_1) have poly-bounded slices (as is readily verified), and it can consequently be verified that $(\text{CI}, \pi_1), (\text{MMC}, \pi_1) \in \text{chopped-co-NP}$ and that $(\text{CMI}, \pi_1) \in \text{chopped-}\Pi_2^P$. It is indeed a general fact that when B is a language of pairs where (B, π_1) has poly-bounded slices, the classes $\text{mixed-}\mathcal{C}$ and $\text{chopped-}\mathcal{C}$ coincide, as made precise by the following theorem.

Theorem 21. *Let \mathcal{C} be a classical complexity class closed under many-one polynomial-time reduction. Let B be a language of pairs such that (B, π_1) has poly-bounded slices. Then, B is in $\text{mixed-}\mathcal{C}$ if and only if (B, π_1) is in $\text{chopped-}\mathcal{C}$.*

We now present a formulation of the notion of reduction used in the CDLS framework (see [2, Definition 2.8]).

Definition 22. *Let A and B be languages of pairs. A mixed reduction from A to B is a triple (f_1, f_2, g) of mappings from $\Sigma^* \times \Sigma^*$ to Σ where f_1 and f_2 are poly-length and computable, and g is polynomial-time computable, such that $(x, y) \in A \Leftrightarrow (f_1(x, \text{un}(|y|)), g(f_2(x, \text{un}(|y|)), y)) \in B$.*

In analogy to Definition 20, here we require that the functions f_1 and f_2 are computable.

As a way of showing hardness, CDLS present mixed-reductions from languages of the form $\{\epsilon\} \times Q^+$ where Q^+ is a classical language that is hard. For example, they present the following reductions.

Theorem 23. *(follows from [2, Proof of Theorem 2.10]) There exists a co-NP-complete problem Q^+ such that there exists a mixed-reduction from $\{\epsilon\} \times Q^+$ to CI.*

Theorem 24. *(follows from [2, Proof of Theorem 3.2]) There exists a Π_2^P -complete problem Q^+ such that there exists a mixed-reduction from $\{\epsilon\} \times Q^+$ to CMI.*

We now present a general theorem showing that exhibiting a reduction from a language of the form $\{\epsilon\} \times Q^+$ yields a hardness result with respect to the classes $\text{chopped-}\mathcal{C}$, made precise as follows.

Theorem 25. *Suppose that A and B are languages of pairs such that there exists a mixed reduction from A to B , and let \mathcal{C} be a classical complexity class. If $A = \{\epsilon\} \times Q^+$ where Q^+ is \mathcal{C} -complete, then (B, π_1) is $\text{chopped-}\mathcal{C}$ -hard.*

Corollary 26. *The problem (CI, π_1) is chopped-co-NP -hard; the problem (CMI, π_1) is $\text{chopped-}\Pi_2^P$ -hard.*

Proof. Follows from Theorems 23, 24 and 25. □

The other way in which CDLS show hardness is by presenting a mixed-reduction from a problem that has poly-bounded slices. For example, they prove the following.

Theorem 27. (follows from [2, Proof of Theorem 3.1]) *There exists a mixed-reduction from CI to MMC.*

We show that this form of reduction can be interpreted as a poly-comp reduction, made precise as follows.

Theorem 28. *Suppose that A and B are languages of pairs such that there exists a mixed reduction from A to B . If (A, π_1) has poly-bounded slices, then (A, π_1) poly-comp reduces to (B, π_1) .*

Corollary 29. *There exists a poly-comp reduction from (CI, π_1) to (MMC, π_1) , and hence (by Corollary 26) the problem (MMC, π_1) is chopped-co-NP-hard.*

Proof. Immediate from Theorems 27 and 28. □

At this point, we can observe that the non-compilability results that CDLS obtain can be obtained in our framework. For example, consider the following. As we have seen (and as stated in Corollaries 26 and 29), the problems (CI, π_1) and (MMC, π_1) are chopped-co-NP-hard. This implies that these two problems are not in chopped-PTIME, unless the PH collapses, via Theorem 19. We can also obtain the non-compilability results in (essentially) the form stated by CDLS: by invoking Theorem 21, it immediately follows that the problems CI and MMC are not in mixed-PTIME, unless the PH collapses. We want to emphasize here that the hardness proofs can be carried out using the notions and concepts of our framework.

6. Completeness and hardness for chopped-NP

In this section, we present completeness and hardness results for the class chopped-NP.

Define HAM-PATH to be the problem of deciding, given an undirected graph G , whether or not G contains a Hamiltonian path; define the parameterization γ so that $\gamma(G)$ is equal to the number of nodes in G . The problems 3-SAT and CIRCUIT-SAT are defined as usual. In the context of 3-SAT, ν is the parameterization that returns, given a formula ϕ , the number of variables that appear in ϕ . In the context of CIRCUIT-SAT, $\mu + \nu$ is the parameterization that returns, given a circuit C , the sum of the number of non-input gates and the number of input gates of C . For each $d \geq 2$, we consider d-HITTING-SET to be the problem where an instance is a pair (H, k) consisting of a number $k \geq 1$ and a hypergraph H where each edge has size less than or equal to d , and one is to decide whether or not H has a hitting set of size less than or equal to k . Note that here, all numbers are represented in unary.

Theorem 30. *The following problems are chopped-NP-complete:*

1. (HAM-PATH, γ)
2. (3-SAT, ν)
3. (CIRCUIT-SAT, $\mu + \nu$)

4. $(d\text{-HITTING-SET}, \pi_2)$, for each $d \geq 2$

Proof. We first discuss the **chopped-NP**-completeness of $(\text{HAM-PATH}, \gamma)$. We prove that $(\text{HAM-PATH}, \gamma)$ and $(\text{HAM-PATH}, \text{len})$ are poly-comp interreducible, which suffices by the **NP**-completeness of **HAM-PATH** and Proposition 16. We argue that $(\text{HAM-PATH}, \gamma)$ poly-comp reduces to $(\text{HAM-PATH}, \text{len})$ as follows. There exists a polynomial p such that the size of an instance G with n vertices can be bounded by $p(n)$. We may thus take g to be the identity map and define $s(\text{un}(n)) = \{\text{un}(0), \text{un}(1), \dots, \text{un}(p(n))\}$. We argue that $(\text{HAM-PATH}, \text{len})$ poly-comp reduces to $(\text{HAM-PATH}, \gamma)$ as follows. Under a standard representation of graphs, we have that (for each graph G) $\gamma(G) \leq |G|$. We may thus take g to be the identity map and define $s(\ell) = \{\text{un}(0), \text{un}(1), \dots, \text{un}(\ell)\}$.

A similar argument can be used to show that $(\text{CIRCUIT-SAT}, \mu + \nu)$ and $(\text{CIRCUIT-SAT}, \text{len})$ are poly-comp interreducible. A similar argument can be used to show that $(3\text{-SAT}, \nu)$ and $(3\text{-SAT}, \text{len})$ are poly-comp interreducible, but note the following. For the reduction from $(3\text{-SAT}, \nu)$ to $(3\text{-SAT}, \text{len})$, to obtain a bound of the form $p(n)$ (where p is a polynomial) on the size of an instance with n variables, instead of defining g to be the identity map, we may define g to be the polynomial-time computable function that reduces duplicate clauses, so that there is indeed a bound of the claimed form, since the number of clauses will then be $O(n^3)$.

The **chopped-NP**-hardness of the problems $(d\text{-HITTING-SET}, \pi_2)$ can be observed as follows. It suffices to prove hardness in the case of $d = 2$, as for higher values of d the problem is more general. To prove hardness in the case of $d = 2$, we take g to be the reduction given by the proof of [12, Theorem 7.44] which is a many-one polynomial-time reduction from 3-SAT to the vertex cover problem; since the value of k in the created instance (G, k) is a polynomial function of the number of variables of the original 3-SAT instance, this g can be completed to be a poly-comp reduction from $(3\text{-SAT}, \nu)$. The containment of $(d\text{-HITTING-SET}, \pi_2)$ in **chopped-NP** can be obtained from the known polynomial kernelization of this problem [8, Section 9.1] and Proposition 10. \square

As a way of witnessing the utility of the presented framework, let us discuss how one of the non-compilability results from a previous paper [5] on the parameterized complexity of model checking can be formulated within this framework. Here, by a *unary signature*, we mean a signature containing only unary relation symbols. Define unary-EP-MC to be the problem of deciding, given a pair (ϕ, \mathbf{B}) consisting of an existential positive sentence and a finite relational structure, each over the same unary signature, whether or not ϕ evaluates to true on \mathbf{B} (see the paper [5] for definitions and background).

Proposition 31. *The parameterized problem $(\text{unary-EP-MC}, \pi_1)$ is chopped-NP-hard.*

Proof. Let h be the reduction given in [5], which is a many-one polynomial-time reduction from the CNF satisfiability problem to unary-EP-MC where an instance having n variables and m clauses is mapped to an instance of the form (S_n^m, \mathbf{B}) , where each S_n^m is a sentence. Let g be the map that, given a 3-SAT formula ϕ , eliminates duplicate clauses from ϕ and then maps the result under h . For a 3-SAT formula ϕ with n variables, it will thus hold that there exists a polynomial $C \in O(n^3)$ such that

$\pi_1(g(\phi)) \in \{S_n^0, S_n^1, \dots, S_n^{C(n)}\}$. If we define $s(n) = \{S_n^0, S_n^1, \dots, S_n^{C(n)}\}$, we thus have that (g, s) is a poly-comp reduction from $(3\text{-SAT}, \nu)$ to $(\text{unary-EP-MC}, \pi_1)$, which yields the result by Theorem 30. \square

Acknowledgements. This work was supported by the Spanish project TIN2013-46181-C2-2-R, by the Basque project GIU12/26, and by the Basque grant UFI11/45.

References

- [1] Simone Bova, Florent Capelli, Stefan Mengel, and Friedrich Slivovsky. Expander cnfs have exponential DNNF size. *CoRR*, abs/1411.1995, 2014.
- [2] Marco Cadoli, Francesco M. Donini, Paolo Liberatore, and Marco Schaerf. Pre-processing of intractable problems. *Information and Computation*, 176(2):89–120, 2002.
- [3] Hubie Chen. A theory of average-case compilability in knowledge representation. In *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9-15, 2003*, pages 455–460, 2003.
- [4] Hubie Chen. Parameterized compilability. In *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30-August 5, 2005*, pages 412–417, 2005.
- [5] Hubie Chen. On the complexity of existential positive queries. *ACM Trans. Comput. Log.*, 15(1), 2014.
- [6] Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *J. Artif. Intell. Res. (JAIR)*, 17:229–264, 2002.
- [7] Wenfei Fan, Floris Geerts, and Frank Neven. Making queries tractable on big data with preprocessing. *PVLDB*, 6(9):685–696, 2013.
- [8] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
- [9] Goran Gogic, Henry A. Kautz, Christos H. Papadimitriou, and Bart Selman. The comparative linguistics of knowledge representation. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, IJCAI 95, Montréal Québec, Canada*, pages 862–869, 1995.
- [10] Paolo Liberatore. The size of MDP factored policies. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence and Fourteenth Conference on Innovative Applications of Artificial Intelligence, July 28 - August 1, 2002, Edmonton, Alberta, Canada.*, pages 267–272, 2002.
- [11] Paolo Liberatore and Marco Schaerf. Compilability of propositional abduction. *ACM Trans. Comput. Log.*, 8(1), 2007.

- [12] Michael Sipser. *Introduction to the Theory of Computation, 3rd edition*. Cengage Learning, 2012.
- [13] Chee-Keng Yap. Some consequences of non-uniform conditions on uniform classes. *Theor. Comput. Sci.*, 26:287–300, 1983.

A. Proofs for Section 3.2

We establish a lemma that will be of aid.

Lemma 32. *Suppose that (Q, κ) and (Q', κ') are parameterized problems; that (g, s) is a poly-comp reduction from (Q, κ) to (Q', κ') ; and, that g' is poly-compilable with respect to κ' and Q'' is a language with $x' \in Q' \Leftrightarrow g'(x') \in Q''$. Then, there exists a function $g^+ : \Sigma^* \rightarrow \Sigma^*$ that is poly-compilable with respect to κ such that $x \in Q \Leftrightarrow g^+(x) \in Q''$.*

If one assumes in addition that κ'' is a parameterization and that $s' : \Sigma^ \rightarrow \wp_{\text{fin}}(\Sigma^*)$ is a function such that (g', s') is a poly-comp reduction from (Q', κ') to (Q'', κ'') , then there exists $s^+ : \Sigma^* \rightarrow \wp_{\text{fin}}(\Sigma^*)$ such that (g^+, s^+) is a poly-comp reduction from (Q, κ) to (Q'', κ'') .*

Proof. Suppose that g is poly-compilable (with respect to κ) via $f : \Sigma^* \rightarrow \Sigma^*$ and $c : \Sigma^* \rightarrow \Sigma^*$; suppose that g' is poly-compilable (with respect to κ') via $f' : \Sigma^* \rightarrow \Sigma^*$ and $c' : \Sigma^* \rightarrow \Sigma^*$.

Define $c^+ : \Sigma^* \rightarrow \Sigma^*$ by $c^+(k) = (c(k), \{(k', c'(k')) \mid k' \in s(k)\})$; it is straightforward to verify that this function is computable and has poly-length. Define $f^+ : \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$ to be the function computed the polynomial-time algorithm that, given a pair (a, x) of strings:

- sets $k = \kappa(x)$;
- views a as a string of the form (d, t) where t is a set of pairs;
- computes $x' = f(d, x)$;
- computes $k' = \kappa'(x')$ (which is in $s(k)$); and then, if there exists d' such that $(k', d') \in t$, computes and outputs $x'' = f'(c'(k'), x')$ (if not, the algorithm outputs the empty string).

Define the function g^+ by $g^+(x) = f^+(c^+(\kappa(x)), x)$.

We have that g^+ is poly-compilable with respect to κ . Let $x \in \Sigma^*$ be any string, and set $k = \kappa(x)$. Given the pair $(c^+(k), x)$, the algorithm for f^+ will compute $x' = f(c(k), x)$, will compute $k' = \kappa'(x')$, and will output $x'' = f'(c'(k'), x')$. We have that $x \in Q$ iff $x' = g(x) \in Q'$ iff $x'' = g'(x') \in Q''$.

Suppose that κ'' and s' are as described in the lemma statement. Define s^+ as $\cup_{k' \in s(k)} s'(k')$. It is straightforward to verify that s^+ has the desired properties. \square

Proof. (Theorem 12) Suppose that (Q, κ) poly-comp reduces to (Q', κ') , and that the problem (Q', κ') is in $\text{poly-comp-}\mathcal{C}$. By definition of $\text{poly-comp-}\mathcal{C}$, there exists a function g' that is poly-compilable with respect to κ' and a language $Q'' \in \mathcal{C}$ such that (for each $x \in \Sigma^*$) $x' \in Q' \Leftrightarrow g'(x') \in Q''$. The function g^+ provided by Lemma 32 witnesses that $(Q, \kappa) \in \text{poly-comp-}\mathcal{C}$. \square

Proof. (Theorem 13) Lemma 32 implies directly that if there exist a poly-comp reduction from (Q, κ) to (Q', κ') and a poly-comp reduction from (Q', κ') to (Q'', κ'') , then there exists a poly-comp reduction from (Q, κ) to (Q'', κ'') . \square

B. Proofs for Section 5

B.1. Proof of Theorem 21

Proof. Suppose that B is in $\text{mixed-}\mathcal{C}$ via f and $B' \in \mathcal{C}$. We have $(x, y) \in B \Leftrightarrow (f(x, \text{un}(|y|)), y) \in B'$. To show that $(B, \pi_1) \in \text{chopped-}\mathcal{C}$, define $g(x, y) = h(c(\pi_1(x, y)), (x, y))$ as follows. Define $c(x) = \{(i, f(x, \text{un}(i))) \mid i = 0, \dots, p(|x|)\}$; define h to be a function computed by a polynomial-time algorithm that, on an input $(d, (x, y))$ where d has the form $\{(i, a_i) \mid i = 0, \dots, p(|x|)\}$ and $|(x, y)| \leq p(|x|)$, outputs $(a_{|y|}, y)$; and otherwise outputs a fixed string $z_N \notin B'$. We claim that (B, π_1) is in $\text{chopped-}\mathcal{C}$ via g .

Let $(x, y) \in \Sigma^* \times \Sigma^*$. If $|(x, y)| > p(|x|)$, then $(x, y) \notin B$ and $g(x, y) = z_N \notin B'$. If $|(x, y)| \leq p(|x|)$, then $g(x, y) = h(c(x), (x, y)) = (f(x, \text{un}(|y|)), y) \in B' \Leftrightarrow (x, y) \in B$. Using the fact that f is poly-length, it is straightforward to verify that there exists a polynomial p' such that, it holds that $|g(x, y)| \leq p'(|\pi_1(x, y)|) = p'(|x|)$.

Now suppose that (B, π_1) is in $\text{chopped-}\mathcal{C}$. Then, there exists a function $g(x, y) = h(c(x), (x, y))$ that is poly-comp with respect to π_1 and a language $D \in \mathcal{C}$ such that $(x, y) \in B \Leftrightarrow g(x, y) \in D \Leftrightarrow h(c(x), (x, y)) \in D$. Define $D' = \{(e, x, y) \mid h(e, (x, y)) \in D\}$. Clearly, D' many-one polynomial-time reduces to D (since h is polynomial-time computable), so by hypothesis, $D' \in \mathcal{C}$. Define c' to be a function such that $c'(x, 1^k) = (c(x), x)$. We have $(x, y) \in B \Leftrightarrow h(c(x), (x, y)) \in D \Leftrightarrow ((c(x), x), y) = (c'(x, \text{un}(|y|)), y) \in D'$. Hence, we have $B \in \text{mixed-}\mathcal{C}$ via c' and D' . \square

B.2. A lemma

We use $\mu : \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$ to denote the parameterization defined by $\mu(x, y) = (x, \text{un}(|y|))$.

Lemma 33. *Suppose that A and B are languages of pairs such that there exists a mixed reduction from A to B . Then, there exists a poly-comp reduction from (A, μ) to (B, π_1) .*

Proof. Let (f_1, f_2, g) be the mixed reduction from (A, μ) to (B, π_1) . We define a poly-comp reduction (h, s) with $h(x, y) = f(c(\mu(x, y)), (x, y))$ defined as follows. Define c to be a computable, poly-length function such that $c(x, 1^k) = (f_1(x, 1^k), f_2(x, 1^k))$; define f so that $f((a_1, a_2), (x, y)) = (a_1, g(a_2, y))$. Define $s(x, 1^k) = \{f_1(x, 1^k)\}$.

Let $(x, y) \in \Sigma^* \times \Sigma^*$. Observe that $h(x, y) = f((f_1(x, \text{un}(|y|)), f_2(x, \text{un}(|y|))), (x, y)) = (f_1(x, \text{un}(|y|)), g(f_2(x, \text{un}(|y|))), y)$. We thus have $(x, y) \in A \Leftrightarrow h(x, y) \in B$. Moreover, we have $\pi_1(h(x, y)) = f_1(x, \text{un}(|y|)) \in s(\mu(x, y))$. \square

B.3. Proof of Theorem 25

Proof. We give a poly-comp reduction from (Q^+, len) to (A, μ) . By Lemma 33 (and Theorem 13), this implies that there exists a poly-comp reduction from (Q^+, len) to (B, π_1) , which suffices by Proposition 16. Define g by $g(x) = (\epsilon, x)$ and s by $s(1^m) = \{\mu(\epsilon, 1^m)\}$. It is straightforward to verify that (g, s) is a poly-comp reduction from (Q^+, len) to (A, μ) . \square

B.4. Proof of Theorem 28

Proof. We give a poly-comp reduction from (A, π_1) to (A, μ) . This suffices by Lemma 33 (and Theorem 13). Fix z_N to be a string not in A . Let p be a polynomial witnessing that (A, π_1) has poly-bounded slices. Define $g : \Sigma^* \times \Sigma^* \rightarrow \Sigma^*$ to be the function where $g(x, y) = (x, y)$ if $|y| \leq p(|x|)$, and $g(x, y) = z_N$ if $|y| > p(|x|)$. Define $s : \Sigma^* \rightarrow \Sigma^*$ by $s(x) = \{(x, \text{un}(0)), (x, \text{un}(1)), \dots, (x, \text{un}(p(|x|)))\} \cup \{\mu(z_N)\}$. It is straightforward to verify that (g, s) is a poly-comp reduction from (A, π_1) to (A, μ) . \square